

UNCLASSIFIED

Defense Technical Information Center Compilation Part Notice

ADP023747

TITLE: Signature Evaluation for Thermal Infrared Countermeasure and IED Detection Systems

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Proceedings of the HPCMP Users Group Conference 2007. High Performance Computing Modernization Program: A Bridge to Future Defense held 18-21 June 2007 in Pittsburgh, Pennsylvania

To order the complete compilation report, use: ADA488707

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP023728 thru ADP023803

UNCLASSIFIED

Signature Evaluation for Thermal Infrared Countermine and IED Detection Systems

John F. Peters, Stacy E. Howington, Jerry Ballard, Jr., and Larry N. Lynch
USACE Engineering Research and Development Center (ERDC), Vicksburg, MS
{john.f.peters, stacy.e.howington, jerry.ballard, larry.n.lynch}@erd.c.usace.army.mil

Abstract

A high-resolution, computational suite has been constructed to produce synthetic thermal imagery of vegetated soil surfaces with landmines or other targets. The imagery is created by coupling models for the ground, vegetation, ray casting, and sensor characteristics to produce realistic thermal infrared simulated imagery. These simulations provide information ranging from simple temperature contrasts to high-resolution images comparable to actual sensor images that can be used to evaluate or train automatic target recognition (ATR) systems. Analyses of the ATR results allow development of recommendations for optimal sensing strategies and additional training to improve ATR performance. The modeling and characterization occurs at the centimeter scale, which requires massively parallel computational resources to meet the demands of the simulation.

The models run simultaneously on a single, parallel, or serial computer and communicate using sockets or files. The soil model is a three-dimensional, spatially adaptive, continuous Galerkin, finite element model that simulates partially-saturated flow and heat transport, coupled to two-dimensional surface water flow. The vegetation model simulates infrared absorption, reflection, and transmission by discretized plant leaves and stems. Ray casting provides boundary conditions for the soil and vegetation thermal models, and produces multi-spectral images of energy reflected and emitted from the synthetic scene. Subsurface phase change, distributed root zone moisture uptake and transpiration, and flow through macro pores and cracks are processes under construction. The parallelization of the individual testbed components is relatively straight-forward. The central difficulty in achieving acceptable performance for the computational testbed in a parallel computing environment is the sequencing of data transfers between components.

Example calculations to be presented include a multi-million element simulation for an arid test site that is only

a few meters in its longest dimension. The models are driven with meteorological data and are built using material property data collected at the field site. Synthetic images produced are compared against those from thermal cameras. A long-term goal of this work is to help build parameter estimation software to infer ground state information (soil moisture and physical property distributions) from airborne imagery.

1. Introduction

A technology shortfall was identified with existing and future countermine systems with respect to detection of land-mines and minefields. Future countermine systems must operate in a wide range of geological and environmental conditions and at higher rates of advance than are currently achievable because of high false alarm rates. Many false alarms are triggered by the expression of rocks, plants, anomalies in surface topography, or material heterogeneity in the image. ATRs used to identify targets in images must be trained to discriminate between these geo-environmental features and real threats. Relying on field testing alone to address this issue is not practical because field testing of airborne sensors is very expensive, many test conditions are not controllable, and the success of post-test identification of false alarms is subject to errors in georegistration and to dynamic conditions on the ground.

A possible solution is to augment the database of ATR training images with synthetic images. Previous attempts to quantify and include geo-environmental effects in synthetic imagery were not process-based but relied on observed 'clutter' signatures. These efforts were largely unsuccessful. The purpose of the work described in this paper is to supplement field-based investigations to improve performance of sensors through a better understanding of the interaction among thermal and hydrologic processes of ground features at the scale of sensor resolution^[5]. A key tool for doing that is the computational testbed, which is a predictive tool to

simulate a sensor's view of a scene. Producing realistic synthetic scenes requires assembling an accurate, physics-based model of processes that affect how energy is altered in a scene as it travels from the source to the sensor.

An early application of the computational testbed's moisture-thermal model^[1] investigated distinctive thermal patterns surrounding surface targets. More recent applications of the testbed have explored the visibility of targets and other features in desert and temperate settings under wet and dry conditions. A tool of this sort also will permit analysis of microhabitat and exploration of the interplay between these fine-scale process and the sensing physics as observed from coarser scales common to remote sensing. Moreover, with a credible, physics-based 'forward' modeling tool to produce sensor imagery, comes an opportunity to exploit sensor images to estimate ground state (soil moisture, soil, or vegetation temperature contrast) and material properties (hydraulic and thermal), even using sensors not designed or optimized for that purpose.

2. Testbed Components

The computational testbed is a collection of interdependent models and utilities. Energy inbound to the scene is tracked from direct solar and down welling longwave sources. Energy is reflected, absorbed, and emitted by vegetation, soils, and other objects in the scene. The reflected and emitted energy is tracked back through the atmosphere and the sensor to produce an infrared image. Interactions among the models range from a 'tight' coupling, normally internal to a single code, to episodic boundary condition swapping, to one-directional file handoff.

A. Ray Caster Model

The testbed is organized around the ray caster model, which directs the flow of energy among the model components. Direct and indirect solar energy is cast to all exposed surfaces and define the incident energy for the vegetation and ground models. Radiant energy emitted from surfaces is cast to neighboring surfaces. Reflected and emitted energy is cast to the sensor. As illustrated in Figure 1, the ultimate product of the ray caster is an ideal high-resolution near-surface image that is sampled by the (presumably lower-resolution) sensor model.

The ray caster carries information about the surface geometry of all entities in the test bed. The vegetation and ground models provide to the ray caster data in the form of three-vertex facets that define all surfaces and the materials that make up those surfaces. The ray caster assembles the facet data from each model component into a master description. Objects, such as leaves, which are

translucent, are modeled by twin facets. The vegetation model depicts leaves by single facets which are visible from both sides. Facets in the ray caster are visible from only one side. Therefore, for translucent materials the ray caster pre-processing program duplicates these facets, but with their opposing side made visible, and identifies facet twins in a separate data array. Thus, when routing energy, a portion of the energy not reflected can be transferred to the twin facet as transmitted energy. Data on direct and indirect solar loading, sun orientation, and down welling thermal loading is obtained from the meteorological data base and assembled in a 'light' file. Also included in this file is the frequency content of the visible and thermal energy. The source files from these data as well as sensor-specific data such as view orientation and frequency bands of interest are provided to the ray caster in a scripting file.

The ray caster differs from those used for visualization, in that the boundary flux must be determined for all surfaces and not just those visible to the sensor. Therefore, the energy is routed from the light sources to all exposed surfaces rather than from the sensor to surfaces to light source as it would be if only visualization was of interest. The basis of the ray caster is a simple visible surface algorithm using Z-buffers (Figure 2). Several Z-buffers are maintained to route the various components of energy. For each light source, energy is compiled into a pixel array upon which the exposed surfaces are projected. Exposure of a surface is determined by the closest surface projected on the pixel. The incident energy is compiled into the facet total which can then be transmitted to the other model components as part of the energy flux boundary condition. Therefore, resolution of the energy input is determined by the facet size.

The ray caster provides the energy flux to the other testbed components as a facet-based list. In turn, each component returns to the ray caster physical temperatures at the vertices of the facets. Reflected energy from the lighting pixel array is projected on a similar pixel array defined by the viewing (sensor) orientation. Energy is compiled into a pixel from a surface if it is exposed to both viewing array and lighting array. In this way, shadows are resolved to the resolution of the pixel size. Added to this energy is thermal energy emitted from the surface as computed from the temperature. The temperature is interpolated to the pixel location from the values at the facet's vertices. Energy transmitted through materials is compiled from twin facets in a similar manner. The viewing pixel array thus receives energy from both thermal and visible sources.

All operations with energy is wavelength dependent. To accommodate this dependency, energy is stored as six integers ranging in value from 0 to 1,024 W/m^2 . Each integer represents the energy within a band. Properties

for reflectance, transmittance, and emissivity are likewise expressed by six components. Energy transfer operations are performed band-wise. Therefore, the ray caster logic contains no wave-length dependence. The physical meaning of each band is controlled by the values contained in the material property and light files. Greater or lesser frequency resolution can be obtained within particular bands depending on these data. The allocation of energy among the six bands depends on the mode of operation. The ray tracer has three modes of operation:

1. Distributes incident energy to all exposed facets in the scene, accounting for shadowing.
2. Distributes radiated heat from the ground to the plant facets.
3. Renders an ideal multi-spectral image that is sampled by the sensor model. This final rendering includes not only reflected energy, as is common of most rendering algorithms, but includes thermal energy radiated from objects. The result of the rendering process is an image file that contains radiance values for six spectral bands.

Thus, for example, the computation of boundary flux (modes 1 and 2) can use energy averaged over a wide range of frequencies, while the final ray casting to the sensor (mode 3) can place resolution in the narrow bands of interest for the sensor being simulated.

B. Moisture/Thermal Model

A key component of the energy balance is heat emitted or absorbed by the soil. For imagery, only the behavior at the ground surface is of interest. But, heat transport must be modeled in three dimensions to capture the surface expression of subsurface heat transport processes. Thermal properties depend on a soil's moisture content. Exploring the effect of hydrologic processes requires that water movement be simulated as well. Moisture and thermal energy movement in the soil is computed with the ADH (Adaptive Hydraulics/Hydrology Model)^[7]. The model includes local refinement and coarsening (h refinement only) of the computational mesh based on error indicators.

Moisture movement through the soil is estimated by solving Richards' equation in three dimensions (Eq. 1) using a traditional, continuous Galerkin, finite element method on tetrahedra.

$$S_s S(\psi) \frac{\partial \psi}{\partial t} + \eta \frac{\partial S(\psi)}{\partial t} = \nabla \cdot [K_s k_r(\psi) \nabla (\psi + z)] + W, \quad (1)$$

where ψ is pressure head, S_s is the specific storage, which accounts for water compressibility and aquifer elasticity, $S(\psi)$ is the water saturation or volumetric fraction of pore space occupied by water, η is the porosity or

volumetric void fraction, K_s is the water-saturated hydraulic conductivity, k_r is the relative permeability of the media, and W is a source/sink term. In this formulation, the z axis is the vertical direction oriented positively upward. Both S and k_r are functions of ψ . K_s and S_s are provided as data. The model also includes surface water equations to approximate runoff, puddle development, ponding, and infiltration. Choices of surface water flow equations include kinematic wave, diffusive wave, and full shallow water. The groundwater and surface water are coupled and exchange fluxes each time step.

An energy balance equation to simulate heat conduction, convection, and surface heat exchange was added to the moisture model of the original ADH. Thermal energy is transported by a standard conduction and convection equation

$$(\rho c)_m \frac{\partial T}{\partial t} + (\rho c)_f \mathbf{v} \cdot \nabla T = \nabla \cdot (K_T \nabla T) + q_m, \quad (2)$$

where ρ is density, c is specific heat, T is temperature, \mathbf{v} is flow velocity, K_T is thermal conductivity, and q_m is a source of energy. Specific heat and thermal conductivity depend on the soil moisture. Mixture coefficients are computed by volumetric averaging of components for thermal conductivity and mass averaging of components for specific heat. Surface heat exchange includes short-wave input, longwave input, longwave emitted, sensible heat, latent heat, and precipitation heat^[4].

C. Vegetation Model

To understand the role of the energy budget components in a canopy-soil system we modeled the vegetation as discrete elements over the soil mesh. The vertical variations of the energy budget for the discrete vegetation elements are described in terms of the net radiation (R_n), sensible (convective) (H), and latent heat (LE), and ground heat flux (G). The energy budget is given as

$$R_n - LE - H - G = 0. \quad (3)$$

Temporal variations of sensible and latent heat fluxes are modeled as a function of vertical position of the discrete vegetation element by linking the latent and sensible heat fluxes to the measured wind speed and relative humidity. The near surface wind and relative humidity profile is based on a model driven by the measured wind and humidity. Solving the energy budget for each element at each time increment provides information on the vertical and temporal variation of vegetation element temperature, latent heat flux, and sensible heat flux. The latent heat modeled as a traditional, continuous Galerkin, finite element method on tetrahedra.

$$LE = \lambda(T_v) \frac{\rho_v^* - h_a \rho_a^*}{r_s + r_a}, \quad (4)$$

is coupled to the wind speed and relative humidity just above the vegetation through to the vegetation discrete element and is dependent on vapor density ρ_v and aerodynamic resistance r_a . The latent heat also depends on the air density at a specified $h_a \rho_a^*$ the saturation vapor density in the canopy ρ_v^* and the element layer modeled stomatal resistance r_s . The stomatal resistance sub-model includes the air temperature, relative humidity, global radiation, and vegetation type. The sensible heat modeled as

$$H = \rho_a C_p \frac{(T_l - T_a)}{R_a}, \quad (5)$$

where T_l is the leaf temperature, ρ_a is the air density, C_p is the specific heat of air at constant pressure, and T_a is air temperature from the measurements made just above the canopy.

D. Sensor Model

The testbed also includes atmospheric and sensor effects on the final image. The sensor model transforms the initial ideal image produced by the ray tracer through a sequence of steps illustrated in Figure 3. The ideal image (contained in a file with extension *rgb*) is "down sampled" to produce an image with resolution consistent with that of the sensor. The near-surface image passes through the industry-standard atmospheric radiance and transmission model, MODTRAN^[2] and a model to reproduce the optical blurring and sampling effects and correlated noise of the specific sensor. The combined effects of blurring are incorporated in the modulation transfer function (MTF). The blurred image is created from a convolution of the down-sampled image with the MTF. The convolution operation is performed in frequency space. The radiance values of the blurred image are then converted to voltages through a process that includes the effects of sensor sensitivity and noise. The final sensor output is converted to a digitized form that is written to a file in tagged image file format (TIFF). All sensor operations are computed band-wise, although, depending on the sensor capabilities, only selected bands are processed through the sensor.

E. Scene Development

While not a specific testbed component, the success of the simulations depends on the realistic depiction of scene geometry. Scene geometries are created from triangular facets, which are effectively triangular finite

elements oriented in three dimensions (Figure 4). Facets are used to form vegetation and other ground objects as well as the ground surface and the boundaries of the ground model. In addition, the volumes within the faceted surfaces are represented by tetrahedral finite elements. The development of quality finite element meshes of heterogeneous ground combined with the various geometries of other scene objects is a formable task and is beyond the scope of this paper. Details of the process can be found in a companion paper^[3].

3. Parallelization of the Testbed

The problems under study present a classical need for fine spatial detail in the simulation while maintaining a large domain size to permit meaningful comparison against field data. Thus, each of the simulation codes was rewritten to function on parallel computer architectures using the Message Passing Interface (MPI). Each of the testbed components (ground model, vegetation model, and ray tracer) had been parallelized individually. In particular, ADH, the basis of the ground model, was intended for use in a high performance computing environment from its inception^[6]. The difficulty in creating an efficient suite of programs was to make the components operate together for time-transient computations.

A. Execution Sequence and Inter-Model Communication

The three components in the ground model suite must exchange information. The ray caster provides inbound thermal energy to the soil model and the vegetation model. These models provide physical temperatures back to the ray caster. These temperatures are coupled with wavelength-dependent emissivity to estimate emitted energy from the scene. Because the codes are continuing to evolve and are each complicated in their own right, the codes were kept separate. The codes communicate at a prescribed time interval using UNIX sockets or files with semaphores. In theory, the codes could be run simultaneously on different computer platforms. Thus far, the codes have been run together only on a single platform.

The flow of information is complicated by the timing of model interaction (Figure 5). The ground and vegetation models are coupled through the ray tracing model (modes 1 and 2). At each time interval, both models must receive information on incident energy flux from the ray tracer in mode 1. During a time interval, the vegetation model must receive radiance from the ray tracer (mode 2) based on temperatures provided by the ground model. At the end of the interval, the ray tracer

(mode 3) renders the image based on temperatures obtained from both models. Each mode of the ray tracer requires a different time to complete its computation cycle. The key difficulty is that each component must wait for information from the other components before proceeding to the next step of the computation. The sequencing of the data flow introduces the potential for significant latency, thus degrading of computational performance. The problem was ameliorated by running three instances of the ray tracer, effectively treating each mode of operation as a separate program. As separate programs, each mode can be sequenced such that it runs as a background process. The one-way data transfer from mode 1 to the models does not require the ray tracer to wait on completion of either model before processing the flux files of the next time step. Therefore, mode 1 can be running continually in the background creating files to be read by the vegetation and ground models. Mode 3 requires input from both vegetation and ground models, but those models do not depend on input from mode 3. The mode 2 computations require temperature data from the ground model and supply flux information to the vegetation model. Both the mode 2 model and the vegetation model run sufficiently fast compared to the ground model and the mode 1 flux calculations that there is generally no wait time involved in the data transfer. The same cannot be said of load balance because of the interdependency of the ray tracing routines and the vegetation model, creating the potential for idle processors.

B. Timings

Initial comparative timings for parallel performance are summarized in Figure 6. The baseline for the comparisons is the limiting case of four processors, whereby each testbed component is run on its own processor. In these particular simulations, the heat up-cast from the ground to the vegetation is not included. The ray caster was parallelized by distributing the inbound light sources. Because there are 17 light sources for most of these simulations, the number of processors allocated for the mode 1 ray caster was set to minimize idle processors. It is notable that the speedup exceeds the ideal up to 16 processors, implying that the computation was memory-limited causing significant loss in efficiency in the baseline four processor configurations. Beyond 16 processors, good performance is obtained.

4. Results

A. Comparison to Field Data

The model suite has been applied to several desert sites for which field data were available. Figure 7 shows the ground surface and plant distribution in a typical synthetic scene. Data included full meteorological data, some temperature data from a column of thermistors, and thermal images from a mid-wave IR camera. Field estimates were made for bulk density, thermal conductivity, albedo, and emissivity.

The model domain was chosen to match the footprint of the boom-mounted sensor cameras. The domain was about 3 meters on a side and 1 meter deep. The spatial resolution on the ground surface was about 12 mm. Resolution was concentrated near the ground surface, decreasing with depth. The resulting mesh was about 2.1 million elements for the soil and about 275,000 elements for the vegetation.

B. Simulated Scenes

Figure 8 shows a sequence of computed images of a three-dimensional domain with color temperature contours. Cooler areas beneath the plants are created by plant shadows. The effect of these shadows propagates into the soil to create thermal shadows. At variance to light shadows, which move in response to change illumination, thermal shadows lag the illumination as a result of the capability of the thermal storage capacity of the ground. The thermal scene presented to a sensor is thus a combination of complex geometries, illumination and shadowing effects, and the time-transient moisture-thermal response of scene components.

C. Example Sensor Images

Simulated sensor images for the example mesh shown in Figure 7 are shown in Figure 9. The image pair allows comparison between day and night scenes. The images illustrate the thermal “cross over” effect that can be seen as a result of relative thermal response the various objects. Plants, which appear cooler during the day relative to the ground, are hotter at night as a result of the faster thermal response of the ground. The thermal shadows beneath the plants are not apparent from the nadir view during the day, but are clearly seen as persistent shadows at night.

D. Component Validation

Individual model components were validated using data from a variety of field measurements. For example,

a ground probe is shown in Figure 10 from which the temperature distributions from a diurnal cycle can be obtained. As can be seen from the accompanying plot, below a depth of 20 cm the temperature remains relatively constant. From this general finding, the domain depth of 1 m appears to amply fall below the active zone. Further, the thermal boundary condition can be defined as a constant temperature approximately equal to the seasonal average temperature. A comparison with the computed data shows that the temperature variation in the upper active zone follows the observed data well. Also note that the small difference between the prescribed boundary condition and the measured temperature had little effect of the temperature profile, and evidently none at all on the variations within the active zone. Similar detailed validation studies have been performed for all testbed components and continue to be performed as new field data are obtained.

E. Image Comparisons

In addition to quantitative comparisons, qualitative comparisons are made between simulated imagery and imagery obtained in field studies. One such comparison is shown in Figure 11 for a scene imaged with a thermal camera. The spectral range of the camera was sensitive to both emitted and reflected energy as is evident from the light shadowing in both images. Overall the images compare well, although the effect of having lower resolution in the simulated image is apparent. The relative brightness within the circular disturbed area in the right-hand side of both images is evident. The shadow of the instrument tripod shown in the thermal camera image is not part of the simulated scene.

5. Application

The focus of this challenge project was the application of the countermine computational testbed to support an airborne field testing program. The testing matrix included an arid and a temperate setting with appropriate soil and vegetation properties. For each setting, four vegetation densities were considered. Each of these scenes was simulated with surface-placed small plastic targets, surface-placed small metal targets, buried large plastic targets, and buried large metal targets. The testing matrix included these 32 distinct scenes. Figure 7 depicts one of the scenes modeled for this application. For each scene, several 5 to 11 day simulations were run for meteorological conditions including dry, dry-to-wet, and wet-to-dry to examine the effects of soil moisture content changes on target detectability. Synthetic sensor images were produced at hourly intervals for each simulation, yielding thousands of images. Each

simulation required about 24 to 72 hours of simulation time on 20 to 32 processors or about 1,500 processor hours per simulation on a Cray XT3. Results of these simulations produced simple guidance on the likelihood of target detection at different times of day and under different meteorological conditions. Analysis of the images suggests a substantial effect of soil moisture content on sensor performance. These simulations have led to a modified schedule for optimum 'hunting' times during and shortly after rain events. Further, when automated target recognition algorithms were run on these synthetic images, false alarms registered on the same features as in the airborne field testing program, thus supporting the idea of using the testbed as a means to train ATRs.

6. Future Model Improvements

Significant moisture movement in the shallow subsurface occurs as vapor phase transport. This need requires the simulation of air movement and precludes the use of Richards' equation. At the time of this writing, a prototype two-phase flow model has been constructed and is being transitioned to the moisture/thermal model. Also under construction is a root zone model to represent plant uptake of water. This addition will couple the soil model and the vegetation model directly. Local variation in wind speed can affect dramatically the surface cooling for both soils and vegetation. Therefore, the current ground model suite is being extended to include a near-surface, three-dimensional air flow model.

7. Conclusion

A computational testbed was built to permit the exploration of fine-scale interaction of hydrologic and thermal processes on vegetated soil surfaces and to supplement field testing of sensor performance. The testbed produces realistic synthetic images that are comparable in resolution and complexity to field-collected images. Initial false alarm analysis of the synthetic and field-collected images shows similar behavior. The testbed components are continually being validated against field data, particularly as data are obtained for new climatic conditions. The simulated imagery is currently being used to predict sensor performance under a variety of operational settings. The experience obtained from these efforts also contributes materially to the fidelity of future testbed simulations. To date, several thousand images have been produced and analyzed, providing data for evaluation of automated target recognition algorithms. The testbed development has been concomitant with field and laboratory investigations from which improvements in the

fundamental physics represented by the testbed models are made.

Acknowledgments

Tests described and the resulting data presented herein, unless otherwise noted, were obtained from research sponsored by the Department of Defense's Institute for Maneuverability and Terrain Physics Simulation and the Countermining Phenomenology Program. Work was conducted by the Engineer Research and Development Center (ERDC). Permission was granted by the Director of the Geotechnical and Structures Laboratory, ERDC to publish this information.

References

1. Albert, M., G. Koh, G. Koenig, S. Howington, J.F. Peters, and A. Trang, "Phenomenology of Dynamic Thermal Signatures Around Surface Mines." *Proceedings of the International Society of Optical Engineering (SPIE) Defense and Security Symposium*, Orlando, FL, March 2005.
2. Berk, A., L.S. Bernstein, and D.C. Robertson, "MODTRAN: A Moderate Resolution Model for LOWTRAN 7." *Air Force Geophysics Laboratory Technical Report GL-TR-89-0122*, Hanscom AFB, MA, 1998.
3. Eslinger, O.J., A. M. Hines, S.E. Howington, J.R. Ballard, J.F. Peters, B. White, and P. McAllister, "Omicron: Rapid Mesh Generation on HPC Platforms for the Study of Near Surface Phenomena with Remote Sensing." *Proceedings of the 2007 High-Performance Computing Users Group Meeting*, Pittsburgh, PA, USA, Washington, DC, DoD HPC Modernization Office, 2007.
4. Frankenstein, S. and G. Koenig, "Fast All-season Soil Strength (FASST)." *Technical Report ERDC/CRREL SR-04-1*, Engineer Research and Development Center, Hanover, NH, 106 pages, 2004.
5. Howington, S.E., J.F. Peters, J. Ballard, T. Berry, L. Lynch, and C. Kees, "A Suite of Models for Producing Synthetic, Small-scale Thermal Imagery of Vegetated Soil Surfaces." *Proceedings of CMWR XVI International Conference for Computational Methods in Water Resources*, Copenhagen, Denmark, edited by Philip J. Binning Peter Engesgaard, Helge Dahle, George F. Pinder and William G. Gray, June 19-22, 2006.
6. Howington, S.E., R.C. Berger, J.P. Hallberg, J.F. Peters, A.K. Stagg, E.W. Jenkins, and C.T. Kelley, "A Model to Simulate the Interaction between Groundwater and Surface Water." *Proceedings of the 1999 High-Performance Computing Users Group Meeting*, Monterey, CA, USA, Washington, DC, DoD HPC Modernization Office, 1999.
7. Schmidt, J.H., *A 3-D Adaptive Finite Element Method for Transport Processes*, Master's Thesis, The University of Texas at Austin, Austin, TX, 1995.

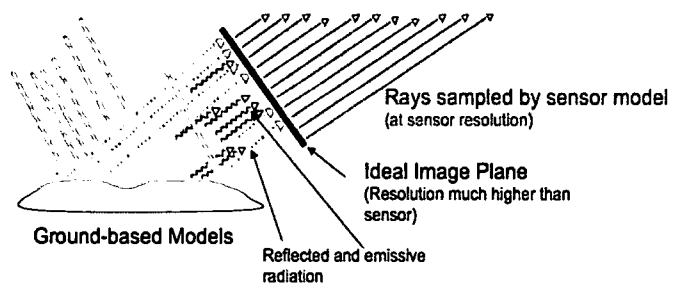


Figure 1. Flow of incident, reflected, and emitted energy through the ray caster to produce high-resolution ideal image to be "sampled" by the sensor model




Figure 2. Use of simple Z-buffer to determine incident energy on exposed surfaces




Figure 3. Data flow diagram for sensor model




Figure 4. Example of a synthetic scene and the underlying meshes




Figure 5. Sequence of data transfers among testbed components during a time step




Figure 6. Speedup obtained with multiple processors (baseline is four processors)




Figure 7. Nadir view of mesh showing relative resolution of ground model size of objects. Relative elevation is indicated by colors.




Figure 8. Temperature variation of a desert scene for selected time in a diurnal cycle (red is the hottest)




Figure 9. Comparison of simulated sensor images for day and night scenes depicting examples of persistent thermal shadows and thermal "cross over", in which cooler objects during the day can appear hotter at night

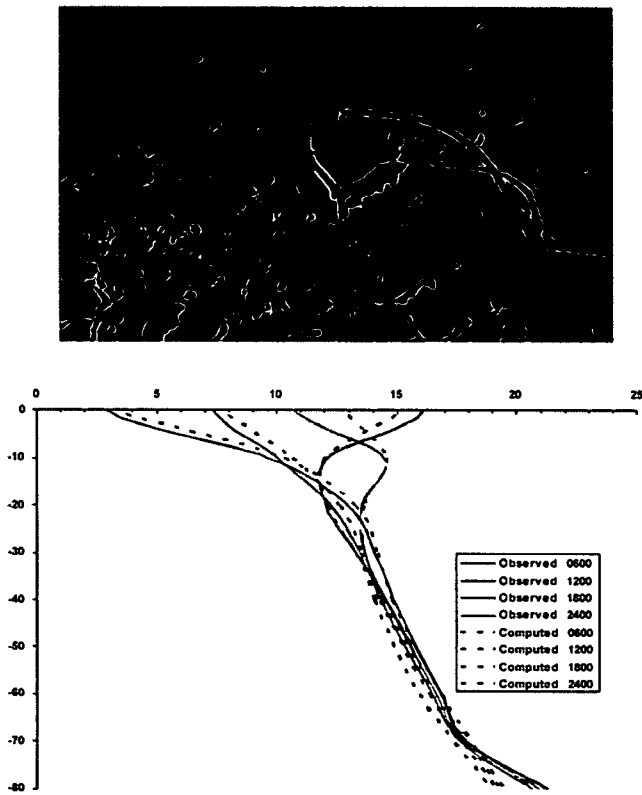


Figure 10. Temperature probe and data for a diurnal cycle showing comparison between observed and computed ground temperatures




Figure 11. Comparison of IR sensor image (left) versus computed image (right). Note that in this case, the resolution of the computed image is less than that of the actual image.